



# UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE

United States Patent and Trademark Office

Address: COMMISSIONER FOR PATENTS

P.O. Box 1450

Alexandria, Virginia 22313-1450

www.uspto.gov

| APPLICATION NO.  | FILING DATE | FIRST NAMED INVENTOR | ATTORNEY DOCKET NO. | CONFIRMATION NO. |
|--|-------------|----------------------|---------------------|------------------|
| 10/706,345   | 11/12/2003  | David Chimitt        | TN313               | 6037             |
| 27276  | 7590        | 09/29/2010           | EXAMINER            |                  |
| UNISYS CORPORATION<br>Office of the General Counsel<br>801 Lakeview Drive, Suite 100<br>MailStop: 2NW<br>BLUE BELL, PA 19422 |             |                      | PARK, ILWOO         |                  |
|  |             |                      | ART UNIT            | PAPER NUMBER     |
|  |             |                      | 2182                |                  |
|  |             |                      | MAIL DATE           | DELIVERY MODE    |
|  |             |                      | 09/20/2010          | PAPER            |

**Please find below and/or attached an Office communication concerning this application or proceeding.**

The time period for reply, if any, is set in the attached communication.



UNITED STATES PATENT AND TRADEMARK OFFICE

Commissioner for Patents  
United States Patent and Trademark Office  
P.O. Box 1450  
Alexandria, VA 22313-1450  
[www.uspto.gov](http://www.uspto.gov)

**BEFORE THE BOARD OF PATENT APPEALS  
AND INTERFERENCES**

Application Number: 10/706,345  
Filing Date: November 12, 2003  
Appellant(s): CHIMITT ET AL.

---

Richard J. Gregson  
For Appellant

**EXAMINER'S ANSWER**

This is in response to the appeal brief filed 7/8/2010 appealing from the Office action mailed 6/9/2009.

**(1) Real Party in Interest**

The examiner has no comment on the statement, or lack of statement, identifying by name the real party in interest in the brief.

**(2) Related Appeals and Interferences**

The following are the related appeals, interferences, and judicial proceedings known to the examiner which may be related to, directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal:

None.

**(3) Status of Claims**

The following is a list of claims that are rejected and pending in the application:

Claims 1-5, 7, 9, 11, 12, 14-16, and 18-25.

**(4) Status of Amendments After Final**

The examiner has no comment on the appellant's statement of the status of amendments after final rejection contained in the brief.

**(5) Summary of Claimed Subject Matter**

The examiner has no comment on the summary of claimed subject matter contained in the brief.

**(6) Grounds of Rejection to be Reviewed on Appeal**

The examiner has no comment on the appellant's statement of the grounds of rejection to be reviewed on appeal. Every ground of rejection set forth in the Office action from which the appeal is taken (as modified by any advisory actions) is being

maintained by the examiner except for the grounds of rejection (if any) listed under the subheading "WITHDRAWN REJECTIONS." New grounds of rejection (if any) are provided under the subheading "NEW GROUNDS OF REJECTION."

#### **(7) Claims Appendix**

The examiner has no comment on the copy of the appealed claims contained in the Appendix to the appellant's brief.

#### **(8) Evidence Relied Upon**

The following is a listing of the evidence (e.g., patents, publications, Official Notice, and admitted prior art) relied upon in the rejection of claims under appeal.

US2003/0158836 A1

VENKATESH et al

8-2003

#### **(9) Grounds of Rejection**

The following ground(s) of rejection are applicable to the appealed claims:

1. **Claims 1-5, 7, 9, 11, 12, 14-16, and 18-25 are rejected under 35 U.S.C. 112, first paragraph, as failing to comply with the written description requirement.**

The claim(s) contains subject matter which was not described in the specification in such a way as to reasonably convey to one skilled in the relevant art that the inventor(s), at the time the application was filed, had possession of the claimed invention. The Specification including original claims does not include "file system." Thus, the Specification does not support "file system" or "before the IRP reaches a file system associated with the IRP."

**2. Claims 1, 2, 4, 5, 7, 9, 11, 12, 14, 16, and 18-25 are rejected under 35 U.S.C. 102(e) as being anticipated by Venkatesh et al. [US 2003/0158836 A1].**

As for claim 1, Venkatesh et al teach a method for processing input/output request packets (IRPs) directed to a Data Volume for providing a single logical storage device to users and applications of a computing system [e.g., "meta file system that appears to a user or application program to be a single file system" in paragraph 0083; "meta file system manager...permitting users and application programs...as if the meta file system were one large conventional file system" in paragraph 0035], the Data Volume having a meta-data extent and at least one data extent [e.g., metadata 158, data 157 stored on cached disk array 114 in fig. 10; or file system metadata, directories, and files in fig. 1], wherein the meta-data extent and at least one data extent are Basic Volumes, and the method is implemented above a Basic Volume Manager [e.g., "a storage layer for cache management and mapping of logical storage locations to physical storage locations on disk" in paragraph 0006; or file systems 119-122 in fig. 8], the method comprising the steps of:

intercepting an initial IRP [e.g., "client makes a request for access to a file stored in a cached disk array 110 over the data network 111", the request is received by NFS FAP software module 141 or CIFS FAP software module 142 in paragraph 0061, fig. 9] before the IRP reaches a file system [e.g., "file systems A:, B:, C:, D: 119-122 on cached disk array 114", file system 121 in figs. 8 and 10] associated with the IRP;

evaluating [e.g., "translations of NFS requests to the intended physical file storage devices" in paragraph 0057; or "These subsequent requests are interpreted by

the NFS routines 141 and forwarded to the meta file system manager 146" in paragraph 0061] the initial IRP by a first volume filter [e.g., "software structure that is replicated in each data mover 115" in paragraph 0057] associated with the meta-data extent to determine the meta-data extent to handle the IRP [e.g., 'meta file system manager 146 includes a routine for accessing the root directory and checking the meta file system flag and a routine for accessing a directory and checking an attribute' in paragraphs 0059, 0060, "meta file system routines to recognize and interpret external links to objects in other file system cells" in paragraph 0059; "the meta file system manager 146 may need data or metadata of a file system object" in paragraph 0063];

directing [e.g., "forwarding the request through the Virtual File System (VFS) to the meta file system manager 146" in paragraph 0062; "The meta file system manager 146 includes a routine for accessing the root directory of a file system cell" in paragraph 0059; "In such case, meta file system manager 146 accesses a file system routing table 148 in order to determine the data mover that owns the file system object" in paragraph 0063] the IRP by the first volume filter to the appropriate meta-data extent; and

redirecting [e.g., "forwarding the request to the data mover that owns the file system" in paragraph 0066; "If the meta file system flag for the file system cell is not set, conventional file manager routines in the file system manager are called to access the file system" in paragraphs 0059, 0045] the IRP from the meta-data extent to a second volume filter [e.g., "If the data mover 115 owns the object, then the object is local, and the meta file system manager 146 accesses data or metadata of the file system through UxFS 144. If the data mover 115 does not own the object, then the object is said to be

remote and the data mover 115 is said to be secondary ... the meta file system manager 146 accesses data or metadata of the file system through a Multiplex File System (MPFS) 152" in paragraphs 0063, 0066] associated with the at least one data extent; and

returning a response [e.g., "obtaining metadata" in paragraph 0078] to the initial IRP from the second volume filter associated with the at least one data extent;

wherein the meta-data extent is a first logical drive and the at least one data extent is a second logical drive [e.g., metadata 158, data 157 in fig. 10; "the file system cell named A: is essentially a repository for metadata of the user-visible file system", "file system cells (B:, C:, and D:) containing user or application files" in paragraph 0055];

the Data Volume appears as a single logical storage device to the users and the applications ["meta file system that appears to a user or application program to be a single file system" in paragraph 0083; "meta file system manager...permitting users and application programs...as if the meta file system were one large conventional file system" in paragraph 0035]; and

the meta-data extent comprises configuration information [e.g., "volume label, ownership, access permission, time stamps, updates, consistency flag, etc." in paragraph 0033] for use in setting up and maintaining the Data Volume.

3. As for claim 2, Venkatesh et al teach the IRP is initiated by an originator [e.g., 'client' in paragraph 0061] of input/output (I/O).
4. As for claim 4, Venkatesh et al teach the meta-data extent is associated with a plurality of data extents [paragraph 0033].

5. As for claims 5, 11, and 24, Venkatesh et al teach the plurality of data extents are located on a plurality of physical disks [e.g., data storage 120-122 in fig. 8].

6. As for claim 7, Venkatesh et al teach creating additional IRPs by the volume filter, each additional IRP being derived from the initiated IRP and related to a single data extent [e.g., file lock request in order to obtain a lock, new metadata, and metadata version number in paragraph 0065, notification of a release of the lock in paragraph 0067].

7. As for claim 9, Venkatesh et al teach a method for storing data across at least one physical disk and presenting the data as a single virtual disk [e.g., "meta file system that appears to a user or application program to be a single file system" in paragraph 0083; "meta file system manager...permitting users and application programs...as if the meta file system were one large conventional file system" in paragraph 0035], comprising the steps of:

intercepting [e. g., "client makes a request for access to a file stored in a cached disk array 110 over the data network 111", the request is received by NFS FAP software module 141 or CIFS FAP software module 142 in paragraph 0061, fig. 9] a first input/output request packet (IRP) from an originator before the IRP reaches a file system [e.g., "file systems A:, B:, C:, D: 119-122 on cached disk array 114", file system 121 in figs. 8 and 10] associated with the IRP;

forwarding [e.g., "forwarding the request through the Virtual File System to the meta file system manager 146" in paragraph 0062] the first IRP to a first volume filter



[e.g., meta file system manager 146; "software structure that is replicated in each data mover 115" in paragraph 0057; fig. 9] associated with the meta-data extent;

creating an additional IRP [e.g., file lock request in order to obtain a lock, new metadata, and metadata version number in paragraph 0065, notification of a release of the lock in paragraph 0067] by the first volume filter for each data extent affected by the first IRP;

transmitting ["forwarding the request to the data mover that owns the file subsystem cell" in paragraph 0066] each additional IRP to a second volume filter ["other data mover having software structure that owns the file system" in paragraphs 0063, 0066] associated with each data extent affected by the first IRP;

allowing each additional IRP to pass through the second volume filter associated with volume filter of each data extent affected by the first IRP; and

returning a response ["obtaining metadata" in paragraph 0078] to the first IRP from each second volume filter associated with the at least one data extent to the originator of I/O.

8. As for claim 12, Venkatesh et al teach the data extents affected by the first IRP are located on separate physical disks [e.g., offline storage medium and online storage medium in paragraph 0077].

9. As for claim 14, Venkatesh et al teach a computer system for providing a single Data Volumes of data storage to users and applications of the computing system, the system comprising:

a plurality of storage clients connected to at least one storage server across a computer network [fig. 8];

a plurality of magnetic disks wherein Data Volumes may be created [fig. 3] and virtually presented to said storage clients, each of Data Volumes having a meta-data extent and at least one data extent [fig. 2], the meta-data extent including a first volume filter [e.g., "software structure that is replicated in each data mover 115" in paragraph 0057; fig. 9] adapted to intercept [e.g., "These subsequent requests are interpreted by the NFS routines 141 and forwarded to the meta file system manager 146" in paragraph 0061] and redirect [e.g., "forwarding the request to the data mover that owns the file system" in paragraph 0066; "If the meta file system flag for the file system cell is not set, conventional file manager routines in the file system manager are called to access the file system" in paragraphs 0059, 0045] input/output request packets (IRPs) received from one of the storage clients, before the IRP is received by an associated file system [e.g., "file systems A:, B:, C:, D: 119-122 on cached disk array 114", file system 121 in figs. 8 and 10], to a second volume filter ["other data mover having software structure that owns the file system" in paragraphs 0063, 0066] associated with the at least one data extent, said first volume filter configured to create an additional IRP [e.g., file lock request in order to obtain a lock, new metadata, and metadata version number in paragraph 0065, notification of a release of the lock in paragraph 0067] for each data extent affected by the IRP; the second volume filter associated with each of the at least one data extent returns a response [e.g., "obtaining metadata" in paragraph 0078] to the IRP, and wherein the first and second volume filters are implemented above a Basic

Volume Manager [e.g., “a storage layer for cache management and mapping of logical storage locations to physical storage locations on disk” in paragraph 0006; or file systems 119-122 in fig. 8]; and

a central management facility [e.g., control station 123 in fig. 8] for controlling the at least one storage server;

wherein the meta-data extent is a first logical drive and the at least one data extent is a second logical drive [e.g., metadata 158, data 157 in fig. 10];

the Data Volume appears as a single logical storage device to the users and the applications [“meta file system that appears to a user or application program to be a single file system” in paragraph 0083; “meta file system manager...permitting users and application programs...as if the meta file system were one large conventional file system” in paragraph 0035]; and

the meta-data extent comprises configuration information [e.g., “volume label, ownership, access permission, time stamps, updates, consistency flag, etc.” in paragraph 0033] for use in setting up and maintaining the Data Volume.

10. As for claim 16, Venkatesh et al teach each storage client is presented with a virtual disk including at least one Data Volume having a meta-data extent and at least one data extent [paragraph 0035; fig. 1].

11. As for claim 18, Venkatesh et al teach the at least one data extent is a plurality of data extents and the IRPs are redirected to the data extents based on which data extents are affected by the IRPs [e.g., pointer of the object for accessing the object in

paragraphs 0072, 0073, sending a file lock request using file system ID in figs. 10 and 13].

12. As for claim 19, Venkatesh et al teach each storage client is presented with a particular Date Volume having a meta-data extent and at least one data extent [paragraph 0035; fig. 1].

13. As for claim 20, Venkatesh et al teach the Date Volume is a simple volume [fig. 1].

14. As for claim 21, Venkatesh et al teach the Date Volume is a spanned volume [fig. 1].

15. As for claims 22 and 25, Venkatesh et al teach the Date Volume includes at least three Basic Volumes and a volume filter is logically disposed above said Basic volumes [figs. 2 and 9].

16. As for claim 23, Venkatesh et al teach a volume filter for redirecting input/output request packets (IRPs) sent from an input/output (I/O) originator, the volume filter comprising:

intercepting means for intercepting [e.g., "client makes a request for access to a file stored in a cached disk array 110 over the data network 111", the request is received by NFS FAP software module 141 or CIFS FAP software module 142 in paragraph 0061, fig. 9; "VFS translated NFS Common File System requests" in paragraph 0057] IRPs sent to a meta-data extent associated with a Basic Volume before the IRP is received by an associated file system [e.g., "file systems A:, B:, C:, D: 119-122 on cached disk array 114", file system 121 in figs. 8 and 10];

evaluating means [e.g., "translations of NFS requests to the intended physical file storage devices" in paragraph 0057; or "These subsequent requests are interpreted by the NFS routines 141 and forwarded to the meta file system manager 146" in paragraph 0061] for evaluating IRPs to determine a meta-data extent to handle the IRP [e.g., 'meta file system manager 146 includes a routine for accessing the root directory and checking the meta file system flag and a routine for accessing a directory and checking an attribute' in paragraphs 0059, 0060, "meta file system routines to recognize and interpret external links to objects in other file system cells" in paragraph 0059; "the meta file system manager 146 may need data or metadata of a file system object" in paragraph 0063]; and

redirecting means [e.g., "forwarding the request to the data mover that owns the file system" in paragraph 0066; "If the meta file system flag for the file system cell is not set, conventional file manager routines in the file system manager are called to access the file system" in paragraphs 0059, 0045] for redirecting the IRPs to at least one data extent associated with at least one other Basic Volume wherein a plurality of data extents are associated [fig. 8; paragraph 0010] with an equal number of Basic Volumes; and

creating means for creating an additional IRP [e.g., file lock request in order to obtain a lock, new metadata, and metadata version number in paragraph 0065, notification of a release of the lock in paragraph 0067] for each data extent affected by a redirected IRP

wherein the meta-data extent is a first logical drive and the at least one data extent is a second logical drive [e.g., metadata 158, data 157 in fig. 10];

the Data Volume appears as a single logical storage device to the users and the applications ["meta file system that appears to a user or application program to be a single file system" in paragraph 0083; "meta file system manager...permitting users and application programs...as if the meta file system were one large conventional file system" in paragraph 0035]; and

the meta-data extent comprises configuration information [e.g., "volume label, ownership, access permission, time stamps, updates, consistency flag, etc." in paragraph 0033] for use in setting up and maintaining the Data Volume.

**17. Claim 3 is rejected under 35 U.S.C. 103(a) as being unpatentable over Venkatesh et al. [US 2003/0158836 A1] in view of well known in the art.**

As for claim 3, Venkatesh et al do not disclose the originator of I/O is a Small Computer Interface Target Mode Driver (SCSITMD); however, a Small Computer Interface Target Mode Driver (SCSITMD) for issuing an I/O request for file access is well known in the art. At the time the invention, one of ordinary skill in the art would have been motivated to include the Small Computer Interface Target Mode Driver (SCSITMD) for issuing an I/O request in order to increase applicability for adapting prevalent SCSI connection for accessing files.

**(10) Response to Argument**

The Examiner summarizes the various points raised by the Appellant and addresses replies individually.

a) *Regarding claim rejections under 35 U.S.C. 112, first paragraph for the new matter “file system” or “before the IRP reaches a file system associated with the IRP”, the Appellant shows, from the specification, “above the Basic Volumes so that IRPs are processed by the volume filter prior to being processed by the Basic Volume Manager” for the support [pages 9 and 10 in Appeal Brief].*

The Examiner respectfully disagrees that the Appellant equates the “Basic Volume Manager” to a “file system.” A volume manager is for managing physical or logical volumes used with disk storages, databases, or a file system. There may be a system including both a file system and a volume manager [e.g., Solalis ZFS, Veritas Storage Foundation Basic may be combined a file system and a volume manager]. Even though the Basic Volume Manager could including a file system, a volume manager is not a file system. Nowhere in the Specification discloses “file” or “file system”. Thus, the specification does not support the limitation “file system” or “before the IRP reaches the Basic Volume Manager not a file system”.

b) *The Appellant alleges that the Examiner interprets the “volume filter” of claim 1 as a “data mover” in Venkatesh because the volume filter, which could be implemented as a driver and would need to communicate to the operating system, would need be implemented outside of the native operating system to intercept an IRP before it is received by associated file system and to implement the volume filter above a Basic Volume Manager while a data mover of Venkatesh being a high-end commodity computer could not be implemented above a Basic Volume Manager, nor could it intercept an input/output request packet before it is received by the associated file*

system. Figure 9 illustrates that the file system (144) is part of the data mover (115). The data mover cannot act independently of the file system in figure 9, and therefore cannot perform the functions of the claimed volume filters [pages 8 and 11 in Appeal Brief]

The Examiner respectfully disagrees. Firstly, the claims don't contain arguing limitations, such as the "driver", "operating system", "outside of the native operating system", etc. Further, the Examiner does not interpret the "volume filter" as a "data mover" of Venkatesh. The Examiner equates the "volume filter" as a "software structure that is replicated in each data mover", such as a meta file system manager 146 in fig. 9. which intercepts an input/output request from a client for accessing a file in a cached array disks having file system and implemented above a Basic Volume Manager [e.g., "a storage layer for cache management and mapping of logical storage locations to physical storage locations on disk" in paragraph 0006 or file systems 119-122 located on a cached disk array 114 as viewed within a single file server 110 in fig. 8: the meta file system manager 146 in a file system layer is above the storage layer]. Venkatesh shows those software structure volume filters, such as meta file system manager 146, Unix based File System (UxFS) 144, Multiplex File System (MPFS) 152, and UxFS 159, MPFS 152, are all implemented within a single file server [figs. 8-10] for processing an input/output request to access a file in a file system stored on array disks, respectively.

c) *The Appellant further alleges that the software implemented in Venkatesh's data mover does not intercept IRPs before they are received by the associated file system [pages 11 and 12 in Appeal Brief]*



Again, the data mover of Venkatesh is not the volume filter of claims. Venkatesh teaches that the volume filter [e.g., meta file system manager 146 in a data mover 115 which is hierarchically above the file system 119-122 in figs. 8-10 for the requests originated from the client] intercepts I/O requests [e.g., "These subsequent requests are interpreted by the NFS routines 41 and forward to the meta file system manager 146" in paragraph 0061; "The CIFS server routines 142 receives CIFS request from the client and forward the request through the Virtual File System (VFS) to the meta file system manager 146" in paragraph 0062] before they are received by the associated file system [e.g., file systems 119-122 on the cached disk array 114 in fig. 8].

**(11) Related Proceeding(s) Appendix**

No decision rendered by a court or the Board is identified by the examiner in the Related Appeals and Interferences section of this examiner's answer.

For the above reasons, it is believed that the rejections should be sustained.

Respectfully submitted,

/Ilwoo Park/  
Primary Examiner, Art Unit 2182

Conferees:

/Tariq Hafiz/  
Supervisory Patent Examiner, Art Unit 2182

/Jason D Cardone/  
Supervisory Patent Examiner, Art Unit 2100